AD-A258 681

(2)

Expert Systems for Maintenance Applications

A short research and development task under DLA900–87–D–0017

Delivery Order #0021

Edward W. Page

Clemson Apparel Research Center

October 1992

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)

2. REPORT DATE

3. REPORT TYPE AND DATES COVERED

4. TITLE AND SUBTITLE

"Expert Systems for Maintenance Applications"

5. FUNDING NUMBERS

DLA900-87-D-0017 (C)
DO 0021

6. AUTHOR(S)

Edward W. Page

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Clemson Apparel Research
500 Lebanon Road
Pendleton, SC 29670

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Logistics Agency
DLA-PRM
Room 4B195 Cameron Station
Alexandria, VA 22304-6100

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Unlimited availability

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

This report focuses upon the potential for a combination of multimedia and expert system technologies to address maintenance problems in apparel manufacturing. There is a recognized shortage of skilled maintenance workers in the apparel industry. The judicious application of multimedia expert systems technology can potentially alleviate the problems caused by the lack of skilled maintenance workers and therefore strengthen the competitive position of the US in apparel manufacturing.

The specific objectives of this project were:

1) to assess the needs of the apparel industry relative to equipment maintenance; and
2) to assess the degree to which interactive video technology can be merged with expert systems to meet industry needs.

Maintenance encompasses two primary functions:

1) diagnosis and replacement of faulty machine parts; and
2) adjustment of machine mechanisms to alter the performance in order to accomodate material of differing characteristics or new product specifications.

14. SUBJECT TERMS

Apparel Manufactruring, Maintenance

15. NUMBER OF PAGES

22

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
unclassified	unclassified	unclassified	unclassified

NSN 7540-01-280-5500

Standard Form 298 (Rev 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Executive Summary

This report focuses upon the potential for a combination of multimedia and expert system technologies to address maintenance problems in apparel manufacturing. There is a recognized shortage of skilled maintenance workers in the apparel industry. The judicious application of multimedia expert systems technology can potentially alleviate the problems caused by the lack of skilled maintenance workers and therefore strengthen the competitive position of the US in apparel manufacturing.

The increasing demand for automation in the apparel industry has exacerbated the problem of equipment maintenance. Apparel manufacturers place major emphasis upon maintenance issues when considering options for plant modernization. The inability of a company to recruit and retain well–qualified maintenance personnel is often cited as the reason that a new sewing and material handling technologies are not introduced into a production line. Equipment manufactures must therefore find ways to simplify maintenance procedures for increasingly complex machines.

Expert systems represent a form of artificial intelligence that can potentially alleviate problems caused by the lack of skilled maintenance workers in the apparel industry. Expert systems technology has already been applied to a wide spectrum of diagnostic tasks ranging from medical diagnosis to equipment maintenance. Because of the text–oriented nature of conventional diagnostic expert systems, they can not be used effectively by persons with limited reading and typing skills. By combining expert systems technology with multimedia technologies (such as CD–ROMs, touch sensitive displays, graphics, animation, full–motion video and voice input/output) it is possible to implement a powerful diagnostic advisor that will allow unskilled maintenance workers to perform at a level approaching that of an experienced technician.

To demonstrate the utility of multimedia expert systems in apparel manufacturing, a prototype system should be implemented. The prototype should be an interactive system incorporating artificial intelligence and multimedia technologies to assist an unskilled maintenance worker in troubleshooting, adjusting and repairing selected equipment. A specific machine should be chosen for the pilot project. The machine chosen should be one which presents significant maintenance problems. Additionally, there must be a readily available source of human expertise in order to construct the knowledge base. The prototype should be tested at an apparel plant in order to verify its effectiveness.

Table of Contents

1.0 Introduction

This report focuses upon the potential for a combination of multimedia and expert system technologies to address maintenance problems in apparel manufacturing. There is a recognized shortage of skilled maintenance workers in the apparel industry. The judicious application of multimedia expert systems technology can potentially alleviate the problems caused by the lack of skilled maintenance workers and therefore strengthen the competitive position of the US in apparel manufacturing.

The specific objectives of this project were:

(1) to assess the needs of the apparel industry relative to equipment maintenance; and

(2) to assess the degree to which interactive video technology can be merged with expert systems to meet industry needs.

Maintenance encompasses two primary functions:

(1) diagnosis and replacement of faulty machine parts; and

(2) adjustment of machine mechanisms to alter the performance in order to accommodate material of differing characteristics or new product specifications.

These functions are different and at the same time closely related. The machine operator typically has the first level of responsibility for maintenance. The operator is responsible for minor, frequently needed machine adjustments. However, a highly-skilled maintenance person is typically required for machine set up and to diagnose major machine malfunctions. The factory personnel make little distinction between the two maintenance functions listed above because they require the same kinds of cognitive skills and are often performed by the same person.

Machine maintenance is a major problem. There is no question regarding the need for increased use of automation in the apparel industry. But high productivity requires a high level of up time for machines. A failure in a crucial machine can disrupt the flow of materials through the plant and have a marked effect on productivity.

One way of achieving a high level of up time is to maintain an adequate level of spare parts on hand. Another is to perform routine preventative maintenance so that impending failures can be caught before they occur. Finally, when a failure does occur, the required repair should be made quickly to minimize the impact on factory productivity. Since repairs must be made quickly, it is important to have highly-skilled maintenance personnel available whenever the factory is operating.

A staff of highly-skilled maintenance personnel is essential to achieving high productivity in manufacturing. The lack of maintenance personnel is creating business opportunities for

third–party maintenance companies in industries such as the computer industry. But for now, apparel plants are having to supply their own maintenance workers. The inability of a company to recruit and retain well–qualified maintenance personnel is often cited as the reason that a new sewing and material handling technologies are not introduced into a production line.

Technology is central to improved maintenance aids. Numerous industries including computer, automotive and aerospace customarily build products featuring built–in tests to facilitate maintenance. In some cases, a technician might plug in a hand–held terminal to perform diagnostic test. In other cases, the machine might be connected to a remote computer which conducts the test. The use of built–in tests is beginning to be used by US apparel equipment manufactures. For example, AMF Reece of Richmond, VA, the last American manufacturer of industrial sewing machines, has incorporated significant maintenance capability into certain items of their product line.

Expert systems featuring optical–disk, touch–sensitive display and voice input/output technologies can be integrated to create powerful, user–friendly maintenance aids. Although many of the ideas contained in this report might be implemented by individual apparel plants, the burden of maintenance ultimately rests with the machine manufacturer. If manufactures do not produce equipment that is reliable and easy to maintain, there will be no market for their products.

2.0 Expert Systems Technology

An expert system has been defined by Professor Edward Feigenbaum of Stanford University, one of the early pioneers of expert systems technology, as "... an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for solution." In essence, an expert system is a computer program that emulates the decision making ability of a human expert in a narrow domain. Typical expert systems use rules and inferencing to solve static problems such as medical diagnosis. An expert system is really just an alternative way of developing software that concentrates on capturing the knowledge necessary to solve a problem rather than on conventional step–by–step programming techniques. This approach to software development, however, offers several important advantages over traditional approaches. An advantage of using rules is that an applications expert who is not skilled in computer programing can review the knowledge base in plain English. Another advantage is that knowledge can be added incrementally by adding, deleting or otherwise altering the rules. Finally, a rule–based system can discern which rules contributed to a conclusion and hence explain how it arrived at a particular conclusion.

During use, an unskilled user can supply facts to the expert system and receive expert advice in response. A diagnostic expert system can guide a user through a series of tests, using inferencing to home in on a problem. Well designed systems are as easy to use as an automated teller. As illustrated in Fig. 1, an expert system consists of two main parts: the knowl-

edge base and an inference engine. The knowledge base contains both facts and rules. The inference engine uses the facts and rules in the knowledge base together with information supplied by the user to reach conclusions. The inference engine searches for rules that can be satisfied by existing facts. When there are no rules that can be fired, the system prompts the user for information. In a diagnostic application, for example, the user may be asked for certain symptoms that describe the condition of the machine in question. By chaining together the appropriate rules in the knowledge base, the cause of the problem can be determined.

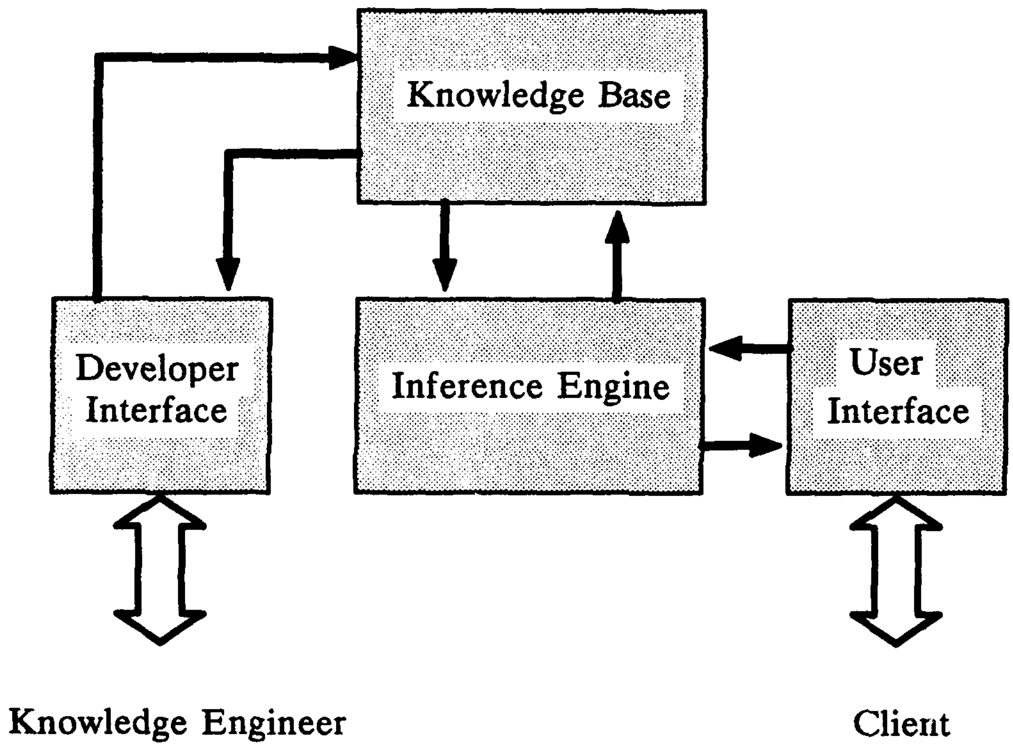

Fig. 1. Components of an Expert System

Rules are usually obtained by interviewing human experts who can solve the problem. As a result, the expert system's knowledge will be limited to a specialized domain. The key idea is that once the expert system is developed, novice users can perform tasks as well as an expert.

3.0 Applying Expert Systems Technology in Apparel Manufacturing

Equipment maintenance is a prime candidate for expert systems technology. Equipment maintenance often requires human intervention to detect, isolate and repair or replace a faulty unit. In most cases, the replaceable unit is a printed circuit board, a power supply or a

cable as opposed to transistors or integrated circuits. While some equipment features built-in tests, a technician is still needed to conduct a series of tests to isolate the fault to a replaceable unit. An expert system has the potential for assisting a novice technician to achieve a level of performance approaching that of an expert.

While there are a number of examples of expert systems that have been successfully applied to the problem of equipment maintenance, today's expert systems are primarily text oriented. The user responds to questions the system poses by typing text on a keyboard and gets recommendations in English sentences on a computer screen. Many of the experienced sewing machine mechanics are unprepared to work with the newest computerized equipment. Moreover, but they often do not have the educational background needed to learn new skills. It is abundantly clear that text-oriented expert systems cannot be used effectively by persons with limited reading and typing skills.

Computers have traditionally been limited to text and still graphics. With recent advancements in VLSI electronics and the emergence of some key standards, computers can now deliver applications that combine text, sound, still graphics, animation, color images and full-motion video images. The power of expert systems technology can be delivered in the form of interactive video, thereby enabling persons with limited reading skills to effectively utilize expert systems.

3.1 Multimedia Technologies

Multimedia refers to the addition of speech input, sound, animation, and video to computer applications. Multimedia systems attempt to emulate the effectiveness of face-to-face communications through an appropriate blend of still and motion video, sound, text and graphics. Multimedia today is a collection of technologies. The challenge is to use the multimedia technologies in a way that can cut costs, communicate more effectively and enhance worker productivity. Presently, the industry is focusing upon multimedia applications in presentations and for training. The use of animation, sound and graphics to convey information helps capture and maintain audience attention.

It is a natural extension of expert systems to begin to encompass multimedia technologies. A multimedia expert system for maintaining equipment in an apparel manufacturing plant, for example, might first guide the repair person through a series of question using graphical displays with a touch-screen input. The computer would communicate with the maintenance technician using speech and graphics. Rather than answering question by typing responses on a keyboard, the maintenance technician would be directed to touch specified points on the screen. This process could be used to collect all information necessary to diagnose a malfunction. Upon completing the diagnosis, the system could make use of both sound and full motion video to give the maintenance technician a short lesson on making the needed repair or adjustment.

Several recent technological innovations can be integrated to produce improved maintenance aids for the apparel industry. This approach would employ an expert system with the latest technology in storage, graphics, interactive displays and voice input/output.

CD–ROM Technology

Compact–disk–read–only–memory(CD–ROM) technology permits the storage and selective retrieval of large volumes of data. Using current technology, a single 4.72 inch disk stores 680 Mbytes of data, the equivalent of 1500 floppy disks or 200,000 pages of text. Of course, CD–ROMs can store animations, audio sequences and full–color images in addition to text. A 256– color full screen image, for example, will require from 200–300 Kbytes. Any portion of the data can be located and displayed within 2 to 3 seconds. The use of data compression techniques can reduce both the amount of memory required to store information as well as the time required to transfer an image to the screen.

A major benefit of CD–ROM technology is that text, graphics, still video images and sound can be stored on the same disk. The cost to produce a video disk master, not including data preparation, is approximately $1000–$2000. Additional copies can be made for approximately $5. a disk.

Inexpensive writable CD–ROM drives promise to facilitate the process of distributing large volumes of data required for multimedia applications. The emergence of low–cost writeabel CD–ROM drives enables the creators of multimedia applications tu bypass the service bureaus that have been needed for the disk mastering process in the past. For small volume quantities, developers can produce CD–ROMs in house for less than $40. per disk. The writable CD–ROM drives needed to produce disks for distribution now cost approximately $8000. However, the standard CD–ROM drive for reading the data cost only about $250.

Video Disk Technology

A modern video disk system allows for the storage of more than 50,000 still pictures or 30 minutes of full motion video per side of a 12 inch disk. Roughly 10 seconds of sound can be recorded in the space required for one still frame. Unlike video tape, the video disk is randomly accessible. This means that frames can be called up and displayed in any sequence. Video disk systems store data in analog form and therefore are not suitable for storing text or computer–generated images. The cost of producing a video disk master is approximately $2000.

Touch–Sensitive Displays

One of the easiest ways to interact with a computer is through a touch–activated display. A touch–screen monitor allows novice users to quickly learn to navigate within a complex software application. The resolution on many of today's systems has become high enough to

permit a user's finger to manipulate a cursor over a single character. Today's touch screen monitors are capable of 1024x1024 touch points. Touch-sensitive displays are well-suited for graphically oriented applications that require the user to make selections or manipulate existing data. Major suppliers are now marketing touch-sensitive products priced between $150 to $500 in volume quantities. To date, most touch-sensitive applications have simply mimicked the function keys on a keyboard; we have only begun to examine novel ways in which to use the technology. In a automated maintenance application, for example, a touch screen monitor may be an ideal solution for interacting with an expert system as it homes in on the faulty component.

The technology places no special demands on software or memory or processor speed. Five competing technologies -capacitive overlay, resistive overlay, surface acoustic wave, piezoelectric and infrared- are found in the commercial marketplace.

Speech Input/Output

More powerful digital signal processing chips, better algorithms and lower integrated circuits prices have caused voice synthesis and recognition products to come of age. A variety of products costing a few hundred dollars are now available for personal computers. Voice synthesis is much less complex than voice recognition but the margin between the two technologies is narrowing as researchers reach a better understanding of the complex nature of speech recognition.

Digital audio techniques can store music and voice. Three sampling frequencies are standard: 44.1 kHz, 22.05 kHz and 11.025 kHz. Higher quality sound requires higher sampling rates and hence additional storage. One minute of low-quality audio requires approximately .66 Mbyte while just over 5 Mbytes of storage is required for one minute of high-quality audio.

Speech Output

Speech output units costing from $200 to $1000 are capable of digitizing human voices, encoding the signals and reproducing speech. The digitized signals may be stored on any convenient medium and selectively accessed under software control. Such units are available with a variety of standard interfaces for personal computers. Text-to-speech devices synthesize speech from text stored in digital form. The quality of the speech is not nearly so good as in digitized speech systems but several commercially available units are quite acceptable. Prices range from $150 to $4000 and at least one system offers more than 99% pronunciation accuracy for more than 20,000 commonly used words.

Speech Recognition

Current speech recognition systems suffer from two major drawbacks:

(1) They are speaker dependent. They must be trained and can recognize only the voice of a particular speaker.

(2) They typically accept only isolated words which means that the speaker must pause between words.

One of the most capable systems produced to date was developed at IBM's Thomas J. Watson Research Center. It recognizes sentences from a 5000 word vocabulary at 95% accuracy. This experimental system requires a large mainframe, three array processors and special purpose hardware. Board-level products appropriate for applications requiring a limited vocabulary of command words are available for personal computers at prices ranging from a few hundred to a few thousand dollars. Such products are capable of recognizing up to approximately 500 words. Systems with a 5000 word vocabulary and 95% recognition accuracy are expected in the marketplace shortly. At the low end, units recognizing in excess of 100 words can be purchased for approximately $300.

Computer Graphics

Expert systems applications place special demands on graphics utilities. High resolution graphics are needed to assist in posing questions to the user and to provide conclusions, explanations and recommendations. Commercially available expert system shells do not always do a good job of permitting developers to incorporate graphics. Ideally, a highly portable, diagnostic expert system needs to incorporate any desired combination of text, graphics and video images on the same monitor.

3.2 Anticipated Benefits

The anticipated impact of employing multimedia expert systems in maintenance applications is increased productivity resulting from widening the range of diagnostic tasks that can be solved by novice technicians. Expert systems applications that currently use only keyboard inputs and textual outputs could benefit from the addition of video images, graphics, an interactive display and voice input/output. The fault isolation techniques should be applicable to a variety of electronic, mechanical, pneumatic or hydraulic systems.

3.3 Components of a Multimedia Expert System

A multimedia system can be developed using a desktop PC. However, a number of interfaces to sound and video peripherals are needed. A development system can easily cost from $10,000 to $30,000; however, once the application is completed, it could execute on low-cost PC's with only minimal additional hardware to support the presentation. A typical development system would employ the following items:

Audio Capture/Playback Board

This board will allow the developer to record, digitize, store and playback sounds. The sounds could come from any variety of sources including speech and music.

Video Image Capture Board

This board is similar to the audio capture board except that it deals with video images. It will allow the developer to capture both still and motion video. The video may come from any NTSC source such as a VCR or a camcorder. With appropriate software, the developer can edit the image including scaling, zooming, panning and cropping.

Video Image Interface

The image interface board accepts still and motion images and displays them in a window of the PC screen. The difficulty has been to display images rapidly enough to give the illusion of motion. If gray-scale video with some motion artifacts are acceptable, the video interface is not required on the delivery system.

Full motion digitized video requires a data rate of 28 Mbytes per second and a one minute segment requires over 1.7 Gbytes of storage. To reduce the amount of storage required for images, an image compression technique is employed. The compression techniques may be applied to both still and full motion images. International standards for image compression have been established. The JPEG standard applies to still images while the MPEG standard was developed for motion video. A video image interface card will provide for compression and expansion of both still and motion video data.

Touch Screen

Unfortunately, many persons given the responsibility for maintenance of apparel manufacturing machines do not have adequate reading skills. Even when reading skills are adequate, the interface needs to be easily adaptable to several different languages in order for the equipment supplier to broaden the market for its products. Touch sens[illegible] displays coupled with graphics oriented user interfaces are therefore well suited for maintenance applications.

Storage

Because of the vast storage requirements for images, a development system often will employ CD-ROM drive. The CD-ROM drive is not necessarily needed on the system that delivers the application.

Computer Graphics Software

The computer graphics software will allow the developer to create text screens and still images and animation for communicating to a technician about how a mechanism such as a looper really works. A wide variety of software for computer graphics is commercially available.

Authoring System Software

The authoring system provides a means of combining graphics, full-motion video, text animation and digital audio into a multimedia presentation. The authoring system must allow the developer to retrieve edit and store data for all media types. Moreover, the authoring system must permit the developer to control the timing and flow of presentations. An authoring system lets a developer create presentations employing text, graphics, animation and video. The developer must be skilled in each medium and must be able to coordinate multiple media into a single coherent system.

Expert System Software

An automated maintenance system used in the apparel industry must be cost effective. It is therefore important to consider the use of low-cost computing platforms and software. This requirement prompted a close look at low-cost expert systems software alternatives that could execute on PC AT platforms.

Regardless of the application, expert systems share a lot of commonality. The user interface and the inference engine need not be altered when a new application is developed. Only the rules and the facts need to be changed in order to implement an expert system for a new application. A variety of expert system shells and programming languages have emerged for facilitating the construction of expert systems.

As a part of the study, we have conducted a rather thorough investigation of CLIPS, a programming language developed by NASA for implementing expert systems. CLIPS is written in the language C and provides high portability across a wide variety of computer systems ranging from PC's to supercomputers. Because it is written in C, it can easily be integrated with other software systems. After writing a number of sample programs in CLIPS, we are convinced that it provides an excellent basis for automating maintenance processes in the apparel industry. An overview of CLIPS is given in Appendix A.

Artificial Neural Networks

Artificial neural networks represent another approach to designing expert systems for maintenance applications. Because neural networks can learn from examples, it is not necessary to interview experts and encode their knowledge in the form of rules. It is possible to train

a neural network using samples of desired behavior. This technique is discussed in Appendix B.

4.0 Conclusions and Recommendations

The increasing demand for automation in the apparel industry has exacerbated the problem of equipment maintenance. Apparel manufacturers place major emphasis upon maintenance issues when considering options for plant modernization. Equipment manufactures must therefore find ways to simplify maintenance procedures for increasingly complex machines.

Expert systems represent a form of artificial intelligence that can potentially alleviate problems caused by the lack of skilled maintenance workers in the apparel industry. Expert systems technology has already been applied to a wide spectrum of diagnostic tasks ranging from medical diagnosis to equipment maintenance. Because of the text–oriented nature of conventional diagnostic expert systems, they can not be used effectively by persons with limited reading and typing skills. By combining expert systems technology with multimedia technologies (such as CD–ROMs, touch sensitive displays, graphics, animation, full–motion video and voice input/output) it is possible to implement a powerful diagnostic advisor that will allow unskilled maintenance workers to perform at a level approaching that of an experienced technician.

To demonstrate the utility of multimedia expert systems in apparel manufacturing, a prototype system should be implemented. The prototype should be an interactive system incorporating artificial intelligence and multimedia technologies to assist an unskilled maintenance worker in troubleshooting, adjusting and repairing selected equipment. A specific machine should be chosen for the pilot project. The machine chosen should be one which presents significant maintenance problems. Additionally, there must be a readily available source of human expertise in order to construct the knowledge base. The prototype should be tested at an apparel plant in order to verify its effectiveness.

Appendix A: An Overview of CLIPS

Introduction

CLIPS is a forward-chaining, rule-based language. The programmer must supply CLIPS with a list of facts and a set of rules. The CLIPS inference engine will then proceed to execute the rules. A fact is just a list of items inclosed in parentheses. For example,

```
(is_broken thread)
```

expresses the fact that the thread is broken. As a second example, the fact

```
(is_connected_to pin_1 pin_23)
```

states that pin 1 is connected to pin 23. Likewise,

```
(test_point_1 10)
```

is a CLIPS fact stating "the voltage at test point 1 is 10."

Rules in an expert system are in the form of a premise followed by a conclusion. For example the rule

```
IF
    the voltage at test point 1 is 10 and
    the voltage at test point 2 is 7  and
    the voltage at test point 3 is 1.2
THEN
    transistor 21 is faulty
```

captures the knowledge relating a specific pattern of voltages at three test points to a particular faulty component. In CLIPS, the above rule would be written in the following form:

```
(defrule RULE_1
    (test_point_1 10.0)
    (test_point-2 7.0 )
    (test_point_3 1.2 )
=>
    (assert (fault_is transistor_21)))
```

The portion of the rule appearing before the "= >" is the premise of the rule and the part appearing after the "= >" is the conclusion. A rule will fire only if its premise is true. In this case, the rule will fire only if there are facts in the facts list that exactly match the premise of the rule. The action of this rule is to place a new fact in the fact list stating that transistor 21 is faulty.

CLIPS provides a powerful facility for matching facts in the premise of rules and for maintaining the fact list. For example, rather than triggering a rule based upon specific facts, the rule can trigger based upon selected fields in the facts appearing in the premise. In this way, only part of a fact needs to be true in order to cause the rule to fire. Additionally, as rules fire, they may not only add facts to the fact list, they can remove facts as well when they are no longer valid. CLIPS supports elementary arithmetic operations, trigonometric functions, conversion functions, string manipulation functions, a variety of predicate functions and functions to control the flow of executions on rule premises. Additionally, CLIPS supports file access. Although CLIPS does not directly implement backward chaining as a part of its inference engine, backward chaining can be emulated using forward chaining rules. Additionally, CLIPS does not directly support a method for dealing with uncertainty. However, a facility for maintaining belief such as the use of certainty factors or fuzzy logic can be added with only moderate effort. Sample programs of this nature were written as a part of our preliminary investigation in order to illustrate that the uses of CLIPS poses no significant obstacles.

Dealing With Uncertainty

Automated maintenance systems must have good facilities for dealing with uncertainty. The following simple example will illustrate some of the problems that are typically encountered and how they may be resolved.

Consider the transistor amplifier of Fig. A–1. The likely faults in resistors and capacitors are either open or short circuits. Likewise, the most likely cause of faults in the transistor are a short or an open between the base and the collector or a short or an open between the base and the emitter. The possibility of a short between the collector and the emitter is also considered. In order to test for a fault, a technician would make voltage measurements at specific test points and would compare the measured voltages to the voltages expected when the circuit was working properly. Table A–1 lists the output voltages expected at the test points for each fault. As shown in Table A–1, the test point voltages provide a "signature" that can be used to identify the fault.

A rule-based system is one approach to the problem of diagnosing faults. Through discussions with an expert technician, a knowledge engineer writes IF–THEN rules that are capable of diagnosing the fault from the symptoms. As shown in Table A–1, certain faults give rise to a unique signature. When C3 is short, for example, the test points uniquely identify the fault. On the other hand, an open capacitor cannot be diagnosed on the basis of the test point voltages since the voltages would read the same as the fault free case. A knowledgeable technician, however, would know to test for an open capacitor by checking to see that the input signal can propagate through the network. Such situations must be recognized and rules must be added to the system to allow such faults to be recognized.

One of the problems with rule-based systems is that they tend to be somewhat brittle. The term brittle means that the system may exhibit rather sharply defined behavior. If a the

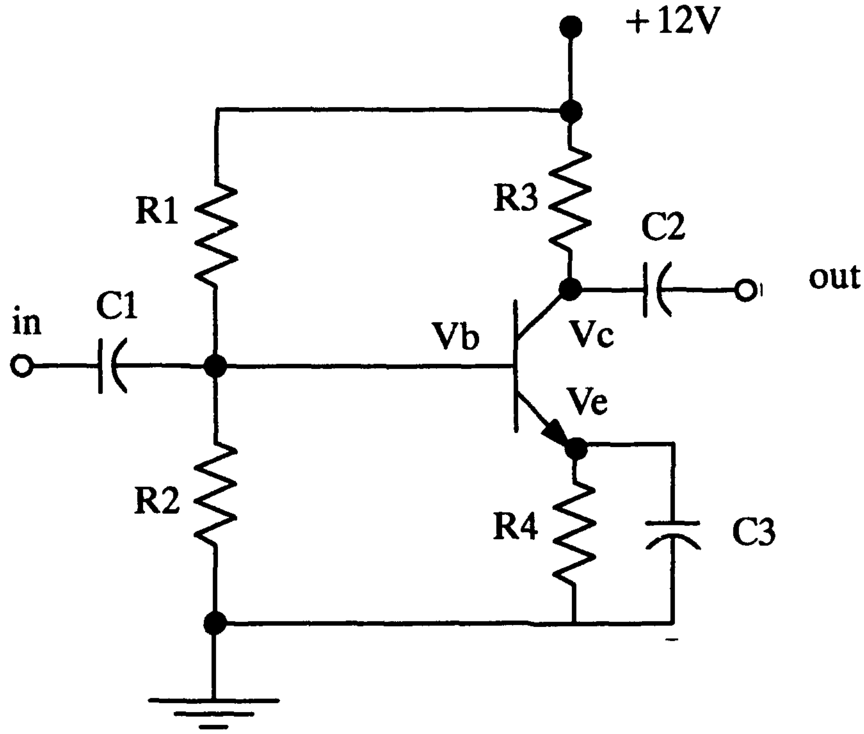

Fig. A-1 Transistor Amplifier Circuit

premise is not matched exactly, the rule will not fire. Yet, an expert human would not necessarily exclude a premise that was nearly matched. For example, we may have a rule as follows for diagnosing faults in the transistor amplifier:

IF the voltage at Test Point 1 = 1.09 and;
  the voltage at Test Point 2 = 0.49 and;
  the voltage at Test Point 3 = 12.00
THEN
  fault = R3_short.

If the voltage at test point 1 was 1.09, the voltage at test point 2 was 0.49, and the voltage at test point three was 11.99, the rule would not fire because the premise is not met exactly. A human, on the other hand, would readily accept 11.99 as being close enough to 12.0 and would conclude that the R3 was faulty. One of the ways of coping with this problem is the use of fuzzy logic.

Fuzzy Logic

According to its inventor, Lotfi Zadeh at the University of California at Berkley, "Fuzzy logic is concerned with modes of reasoning which are approximate rather than exact." People

Table A–1. Fault Signatures for the Amplifier Circuit in Fig. A–1.

TEST POINTS V_b	V_e	V_c	FAULT
1.09	0.49	7.14	none
0.00	0.00	12.00	R1 open
12.00	11.40	Vc < Vb	R1 short
1.70	1.10	Vc < Vb	R2 open
0.00	0.00	12.00	R2 short
0.62	0.02	Vc < Vb	R3 open
1.09	0.49	12.00	R3 short
1.09	0.00	12.00	R4 open
1.09	0.49	7.14	C1 open
1.09	0.49	7.14	C2 open
1.09	0.49	7.14	C3 open
0.70	0.00	Vc < Vb	C3 short
0.62	0.02	12.00	c–b open
1.64	1.04	1.64	c–b short
1.09	0.00	12.00	e–b open
0.05	0.05	12.00	e–b short
1.09	1.09	1.09	e–c short

use fuzzy concepts daily. The concept of "normal height" for a person is not clear cut, yet we all have a concept of normal height. We can capture a concept such as normal height through the use of a membership function. A membership function measures the degree of membership in a fuzzy set. For every element of a fuzzy set, there is an associated membership function. Another way of viewing a membership function is as a representation of the degree to which an object possesses an attribute. For example, the bell–shaped curve in Fig. A–2 can be used to formalize the concept of normal height. A person who is 5'–9" is given a membership grade of 1. Persons whose height is greater or less than 5'–9" have a membership grade that decreases as the height moves away from the center of the bell–shaped curve. In a similar manner, we can capture concepts such as "hot", "higher" and "thick." It is no accident that the first commercial uses of fuzzy logic are in industrial control applications, areas where people have traditionally served as controllers.

Consider now an application is machine maintenance. Suppose that a technician knows that the voltage at a particular test point should read approximately 10 volts when the machine is operating properly. If the technician finds that the voltage is say 9.8 volts, he will probably conclude that the voltage is within acceptable limits. A computer program that has been designed to accept a range of voltages such as 9.8–10.2 volts would reject 9.79 volts. A technician would likely consider 9.79 volts as being too close to reject. The idea of crisp cut-off points is clearly unnatural for a maintenance application. We can capture the concept of "acceptable voltage at test point" using the bell-shaped curve of Fig. A-2. The curve represents the degree of belief that a particular reading is acceptable. A reading of 10 volts is acceptable with the maximum confidence level of 1 while a reading of 9.79 volts has a confidence level of only 0.85.

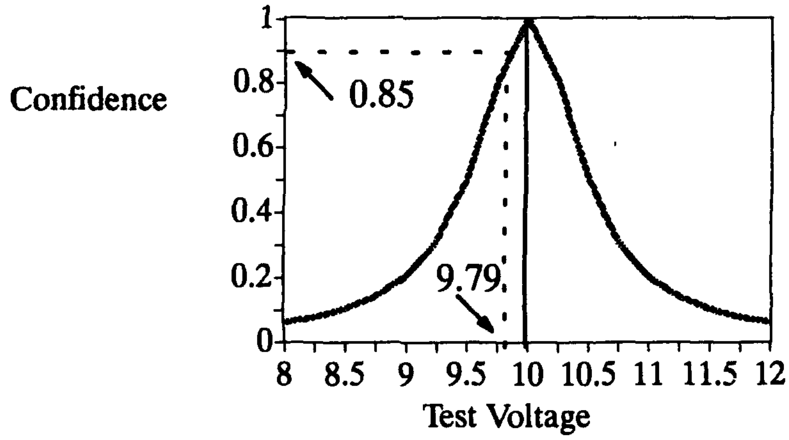

Fig. A-2. Confidence That a Test Point Voltage is 10 Volts

The confidence curve shown in Fig. A-2 was generated by the function

$$f(x) = \frac{1}{1 + \left(\frac{x - \alpha}{\beta}\right)^2}$$

with $\alpha = 10$ and $\beta = 0.5$. This curve represents the fuzzy concept "x is close to α." The parameter β is referred to as the "half-width" since $f(x) = 1/2$ when $x = \alpha \pm \beta$. Note that Eq. 1 always produces a response between 0 and 1

CLIPS does not provide built-in support for uncertainty management. As a part or this study, we implemented two types of uncertainty management to verify that no major problems were posed by the lack of uncertainty management in CLIPS. The two uncertainty management techniques implemented as part of our evaluation of CLIPS are:

(1) the use of certainty factors, as found in several of the commercially available expert system shells; and

(2) fuzzy logic.

Both approaches were implemented in CLIPS and tested; no difficulties were encountered with either method. Because of the diverse forms of reasoning needed for diagnostic systems, both of the above uncertainty management techniques are recommended for a system designed for maintenance applications.

Appendix B: Neural Information Processing

Conventional expert systems are based upon rules. To compile the rules necessary to solve a problem such as machine maintenance, the expert system developer works closely with a domain expert, it this case an expert maintenance technician, to reduce his problem solving expertise to a set of IF–THEN rules. Expert systems have been successful in capturing the expertise of human experts in a wide variety of applications ranging from medical diagnosis to oil exploration. An expert system represents an excellent approach to capturing knowledge that can be verbalized and therefore transferred to a machine. An expert system is not useful in situations in which the knowledge necessary to solve the problem does not exist or can not be verbalized. When a new machine is first placed into operation, there is little practical knowledge available. Additionally, expert maintenance personnel may often base their diagnoses on "patterns" of machine behavior which are difficult to verbalize and express as a set of rules. In such cases it is possible to employ learning algorithms that can learn to diagnose faults from a set of example fault signatures. One approach to learning to diagnose faults from examples is the use of artificial neural networks.

A neural network is a highly–interconnected network of analog processing elements that mimic biological neurons. Current research in the area of neural networks tends to be directed toward the development of models to increase understanding of how the brain works or to construct useful computational devices that are inspired by biological nervous systems. The efficacy of a neural network results from both the analog nature of the computational elements and the high degree of connectivity among the neurons. One of the most important properties of the neural models is that they are capable of working with noisy or incomplete information. A second feature is that they are capable of learning from their environment. Although current neural networks are highly–simplified models of biological nervous systems, they possess computational properties that are applicable to a variety of problems including machine vision, speech processing, pattern recognition, cognitive information processing and combinatorial optimization.

Figure B–1 illustrates a neural network consisting of three layers. Nodes in each layer are connected to nodes in the following layer through adjustable weights that correspond to biological synapses. The first layer serves only to distribute a weighted version of the inputs to the neurons in the second layer. Neurons in the second layer, commonly called the hidden layer, respond to the input signals and propagate their response to the output layer which computes the networks response to the input signal. Such networks can be trained, by adapting the strength of the interconnection, to map input patterns into desired outputs. The neurons in the middle layers are hidden in the sense that their inputs and outputs lie within the network. Each of the signal paths between layers has an associated weight.

For problems that can be cast as classification problems, there is a powerful learning algorithm known as backpropogation that can be used to establish the weights of connections among neurons. In such networks, the representation of knowledge is distributed over many

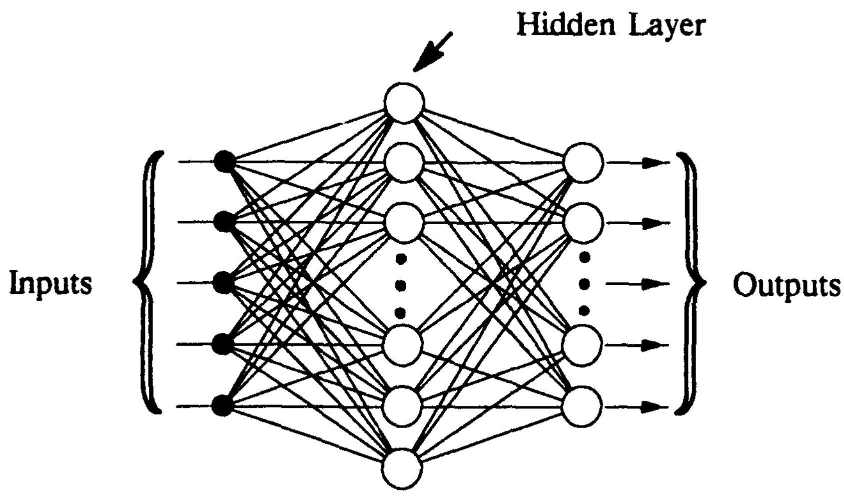

Fig. B–1. Feedforward Neural Network

units. The backpropogation algorithm permits learning to occur solely by revision of connection strengths between neurons. Algorithms such as backpropagation employ samples of desired behavior to adaptively adjust the weights in a manner that causes the network to learn to associate desired output patterns with specific input patterns. Once trained, such networks are typically capable of generalizing their response to provide the correct output for input patterns that were not included in the training set.

To illustrate the utility of neural networks, a neural network was trained on the fault signatures in Table A–1. In addition to achieving 100% accuracy on the training data, this network also exhibited the property of fuzzy systems to correctly classify a fault when the symptoms were reasonably close to the nominal voltage values listed in Table A–1. This project has identified a novel approach to encoding analog inputs to neural networks that can significantly reduce training time. The encoding technique, which we refer to as ensemble encoding, was used to train a network for fault diagnosis that was demonstrated to industry during an "open house" at the Clemson Apparel Research Center in the fall of 1991.